arXiv:1701.00642v1 [cs.AI] 3 Jan 2017

Finding Risk-Averse Shortest Path with Time-dependent Stochastic Costs

Dajian Li* Paul Weng† Orkun Karabasoglu‡

January 4, 2017

Abstract

In this paper, we tackle the problem of risk-averse route planning in a transportation network with time-dependent and stochastic costs. To solve this problem, we propose an adaptation of the A* algorithm that accommodates any risk measure or decision criterion that is monotonic with first-order stochastic dominance. We also present a case study of our algorithm on the Manhattan, NYC, transportation network.

Keywords: Route Planning, Shortest Path, Risk-averse decision-making, Conditional Value-At-Risk, Time-dependent Stochastic Costs

1 Introduction

Shortest path problems have been extensively studied as they are canonical problems that appear in many domains, for instance transportation [2, 7], artificial intelligence [28] or circuit design [25] to cite a few. The standard version of this problem can easily be solved with classic shortest path algorithms such as the Dijkstra algorithm [8] or the A* algorithm [16].

In this paper, we focus more particularly on route planning in transportation networks. While classically route planning operates with deterministic information (e.g., expected travel duration), with the advent of intelligent transportation systems that provide real-time and historical traffic data, it becomes possible to design route planning approaches that take into account the stochastic and time-dependent nature of traffic condition. Indeed, as more and more cities open the access to historical traffic data, it is now possible to estimate a probability distribution over durations for each street at different times of the day. Such information can then serve as input to determine "shortest" paths that takes into account the variability of durations.

More specifically, in this paper we focus on building a risk-averse route planning system for drivers in networks with stochastic and time-dependent costs. For a given origin and destination positions, it determines a shortest risk-averse path with respect to a pre-specified risk measure or decision criterion. With this system, a driver could not only plan their trip in advance, but also avoid possible congestions. Consequently, this system could help reduce in particular travel time, traffic congestion and as a consequence exhaust emissions.

*dajianl@andrew.cmu.edu
†paweng@cmu.edu, http://weng.fr
‡karabasoglu@cmu.edu, http://karabasoglu.com

The contributions of this paper are twofold. First, we propose an adaption of the A* algorithm, which extends and unifies previous algorithms [6, 24] for computing a risk-averse shortest path in transportation networks where costs are stochastic and time-dependent. Our approach can accommodate any risk measure or decision criterion that is monotonic with respect to first-order stochastic dominance. Second, we demonstrate our proposition in the Manhattan, NYC transportation network with Conditional Value-at-Risk as a risk measure.

The paper is structured as follows. The next section discusses the related work. Section 3 recalls the standard shortest path problem and the A* algorithm. Section 4 defines the time-dependent stochastic-cost shortest path problem tackled in this paper. Section 5 presents an adapted version of the A* algorithm to solve our problem. Section 6 demonstrates our solution algorithm to route planning in Manhattan, NYC. Finally, we conclude in Section 7.

2 Related Work

Over the past decades, much effort has been devoted to the solution of the shortest path problem and its many variants.

Classic shortest-path algorithms such as the Bellman-Ford algorithm [4, 11, 19], the Dijkstra algorithm [8] or the A* algorithm [16] have been proposed before 1970s to solve the static version of the problem where edge costs are scalar and constant. However, in route planning, drivers usually value more travel times than distances, which has several implications. Edge costs are generally non-stationary, that is they are a function of time (e.g., driving the same street during peak hours or during normal hours lead to different durations). They also tend to be random, depending on traffic conditions and other drivers. For these reasons, those classic algorithms for static shortest-path problems need to be adapted to this more general setting.

On the one hand, many studies have considered the non-stationary case, i.e., time-dependent shortest path problem (TDSPP). Dreyfus [9] extended the Dijkstra algorithm to TDSPP and Goldberg et al. [15] solved TDSPP with a variant of the A* algorithm. TDSPP has been proven to be solvable in polynomial time under the First In First Out (FIFO) property (i.e., which forbids an earlier arrival time while traversing an edge at a later time) [18] while it reveals to be an NP-hard problem without the FIFO property [23].

On the other hand, since Frank [12] studied stochastic-cost shortest path problem (SSPP), extensive work has been done on this problem (e.g., [31, 21, 14, 24]). Bertsekas et al. [5] considered an even more general class of stochastic shortest path problems (where node transitions are stochastic) and modeled them as a Markov Decision Problem [26]. In SSPP, some researchers, such as Nie et al. [21], aimed at determining a shortest path guaranteeing a given probability of arriving on time. Recently, Gavriel et al. [14] and Parmentier et al. [24] investigated risk-averse versions of SSPP by considering different risk measures, such as conditional Value-at-Risk (CVaR) [10]. However, neither of them considered the case where costs are time-dependent.

Besides, some work also tackles problems where edge costs are both non-stationary and random. Fu et al. [13] considered the problem of expected shortest paths in dynamic and stochastic traffic networks. Chen et al. [6] studied time-dependent stochastic shortest path problems and proposed an adapted A* algorithm with first-order stochastic dominance in order to compute a reliable

shortest path. However, the risk measure they use is Value-at-Risk (VaR) [17], which may not always be the most suitable measure. In this paper, we propose a practical algorithm for any risk measure that is monotonic with respect to first-order stochastic dominance (as in [24]) and test it with CVaR, which may be considered a better criterion than VaR as it has better properties [1] and takes into account not only VaR but also the tail distribution.

3 Background

We first recall the definition of the classic shortest path problem. Let $G = (V, E)$ be a directed graph (e.g., corresponding to a transportation network) where V is a set of nodes (e.g., intersections and landmarks in a city) and $E \subset V^2$ is a set of directed edges (e.g., lanes of streets). The set of successors of a node n is denoted $E^+(n)$, i.e., $E^+(n) = \{n' \in V \mid (n, n') \in E\}$. A path π of length k in G is a sequence of k edges in E: $(n_1, n_2), (n_2, n_3), \ldots, (n_k, n_{k+1})$. For convenience, we write $\pi = (n_1, n_2, \ldots, n_{k+1})$. A subpath of a path is a consecutive subsequence of edges of that path.

The edges of graph G are assumed to be valued by a cost function $c : E \to \mathbb{R}$ (e.g., representing the distance or duration of travel in an edge). We assume that costs are non-negative. By extension, the cost of a (sub)path π, denoted $c(\pi)$, is defined as the sum of the costs of the edges in that (sub)path. Let π_{on} be a subpath from node o to node n and π_{nd} a subpath from node n to node d. We denote $\pi_{on} \oplus \pi_{nd}$ the path obtained from the concatenation of the two subpaths. Obviously, $c(\pi_{on} \oplus \pi_{nd}) = c(\pi_{on}) + c(\pi_{nd})$.

Let $o \in V$ (resp. $d \in V$) be an origin (resp. destination) node. The shortest path problem consists in searching for the path starting from node o and ending in node d that has the lowest cost. Many efficient algorithms, such as the Ford-Bellman algorithm [11, 4] or the Dijkstra algorithm [8], have been proposed to solve this problem. In the case of transportation networks where the number of nodes may be large, those algorithms, even though polynomial in the size of graph G may become impractical. In that case, the A* algorithm may help to determine a shortest path faster.

The A* algorithm, proposed by Hart et al. [16], has been widely accepted as an efficient algorithm to solve the shortest path problem. As it is well-known, we only recall its principle and not its pseudo-code for space reasons. In this algorithm, an extra heuristic information is assumed to be given: for any node n, an estimation $h(n)$ of the cost of the shortest path from node n to destination node d is available. For instance, in transportation networks, where path distances are minimized, $h(n)$ can be defined as the Euclidean distance from node n to destination node d. The A* algorithm finds a path from origin node o to destination node d by exploring a tree of (sub)paths following a best-first-search strategy. In order to choose the best subpath to extend, the A* algorithm usually maintains a priority queue $\mathcal{O}$ of nodes representing subpaths ending in those nodes. The priority $f(n)$ of a subpath π_{on} ending in a node n is defined as the sum of the cost cumulated so far and the heuristic estimation, i.e., $f(n) = g(n) + h(n)$ where $f(n)$ represents an estimation of the cost of of a path to node d whose subpath is π_{on}, $g(n)$ is the cost of π_{on} and $h(n)$ is the heuristic estimation of the cost of a subpath from node n to node d.

Heuristic function $h(n)$ plays a significant role in the A* algorithm by influencing the number of (sub)paths A* algorithm will examine. Besides, whether

the A* algorithm can eventually find the shortest path in the graph depends on the selection of the heuristic function $h(n)$. In order to guarantee the soundness of the A* algorithm, $h(n)$ should satisfy the following inequality: $\forall n \in V$,

$$h(n) \leq \min_{\pi_{nd}} c(\pi_{nd}) \tag{1}$$

where π_{nd} represents a subpath from node n to node d. This property means that the heuristic information provided by $h(n)$ is a lower bound to the best possible cost to reach node d from node n. For instance, the heuristic function defined as the Euclidean distance is admissible. A heuristic function that satisfies inequality (1) is called an *admissible* heuristic function.

4 Problem Statement

We start with some notations. For any random variable X, we denote P_X its probability densition function (pdf), F_X its cumulative distribution (i.e., $F_X(c) = \int_{-\infty}^{c} P_X(x)dx$) and F_X^{-1} the (pseudo)inverse of F_X (i.e., $F_X^{-1}(\alpha) = \inf\{c \in \mathbb{R} \,|\, F_X(c) \geq \alpha\}$).

In a real transportation network, the duration for traversing an edge (i.e., portion of a street) is stochastic and dynamic. Such a network can be represented as a directed graph $G = (V, E)$ as before, however, edge costs are now time-dependent real random variables. For an edge (n, n'), random variable $C_t(n, n')$ denotes the random cost of traversing that edge at time t. We assume random costs take non-negative values (representing durations) and S-FIFO[1] (Stochastic FIFO) [21], which is a natural property in transportation networks, holds.

For a path $\pi = (n_1, n_2, \ldots, n_{k+1})$, its cost $C_t(\pi)$ for a departure time t is also a random variable defined as the sum of the random costs of its edges. It can be written recursively as follows:

$$C_t(\pi) = C_t(\pi') + C_{t+C_t(\pi')}(n_k, n_{k+1}) \tag{2}$$

where $\pi' = (n_1, n_2, \ldots, n_k)$. In a similar fashion, the pdf of $C_t(\pi)$ can be written:

$$P_{C_t(\pi)}(c) = \int_{-\infty}^{+\infty} P_{C_t(\pi')}(x) P_{C_{t+x}(n_k, n_{k+1})}(c - x)dx \tag{3}$$

The problem we tackle in this paper can then be formulated: given a risk-averse criterion or risk measure $\rho : \mathcal{X} \to \mathbb{R}$ (with $\mathcal{X}$ the set of real random variables), we search for the ρ-minimum path π^* for a departure time t, i.e.,

$$\rho(C_t(\pi^*)) = \min_{\pi} \rho(C_t(\pi)) \tag{4}$$

We call π^* a risk-averse shortest path. We assume that criterion ρ satisfies a consistency property that relates ρ to the first-order stochastic dominance, which is a partial order defined over probability distributions [30].

Definition 4.1 First-order stochastic dominance *(FSD) is defined as follows: Let F_1, F_2 be two cumulative distributions, F_1 (weakly) first-order-stochastically dominates (or FSD-dominates) F_2, denoted $F_1 \succsim_{FSD} F_2$, iff $\forall x, F_1(x) \leq F_2(x)$. An illustration of FSD is shown in Figure 1.*

[1] The SFIFO property states that for any confidence level α, leaving later cannot lead to an earlier arrival time: $t \leq t' \implies t + F^{-1}{}_{C_t}(\alpha) \leq t' + F^{-1}{}_{C_{t'}}(\alpha)$ where t, t' are departure times, $C_t, C_{t'}$ random costs of an edge and $\alpha \in [0, 1]$.

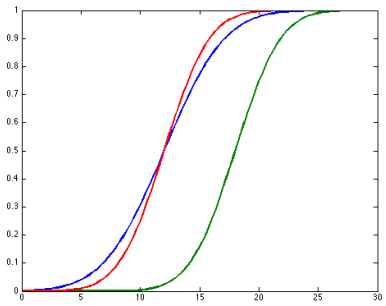

Figure 1: Illustration of First-Order Stochastic Dominance: the green cumulative distributions FSD-dominates the blue and red ones, while the latter two are incomparable.

The consistency property that we assume states that ρ is monotonic with respect to first-order stochastic dominance:

FSD $F_X \succsim_{FSD} F_Y \Rightarrow \rho(X) \geq \rho(Y)$

where X and Y are two real random variables and F_X and F_Y are their respective cumulative distributions. This property is important because it will allow us to prune in the adapted A* algorithm.

Let us introduce another property that states that ρ is increasing with the addition of a non-negative random variable:

INC $\rho(X) \leq \rho(X + C)$

where X and C are two real random variables and C takes non-negative values. In our setting, this is a natural property as random variables represent durations.

To prove that FSD implies INC, we first introduce a lemma[2]:

Lemma 4.1 *Let X be a real random variable and C be a non-negative real random variable. Then, $F_{X+C} \succsim_{FSD} F_X$.*

Then, as a direct consequence of Lemma 4.1, we obtain:

Proposition 4.2 *If ρ satisfies FSD, ρ also satisfies INC.*

As illustrations of ρ, we present three examples, Value-at-Risk, Conditional Value-at-Risk and Expected Utility, which all satisfy FSD (and therefore INC).

Example 1. *Value-at-Risk* (VaR) [17] is a widely-used risk measure in finance. For a fixed $\alpha \in [0, 1]$, it represents the threshold loss value, such that the probability the loss on an investment exceeds this value is α. Formally, in our context, it is defined by: $VaR_\alpha(X) = F_X^{-1}(\alpha) = \inf\{x \in \mathbb{R} \mid F_X(x) \geq \alpha\}$ In other terms, $VaR_\alpha(X)$ is defined in our context as the α-quantile of random variable X. It is well-known that VaR satisfies FSD [3].

Example 2. *Conditional Value-at-Risk* (CVaR) [10], also called *Expected Shortfall* is a risk measure that refines VaR. Because VaR is a threshold value (for a single fixed probability α), it neglects the risk at the tail of the distribution. CVaR remedies this shortcoming of VaR by measuring the expected loss at the

[2]For space reasons, we do not include the proofs.

tail above VaR. CVaR is mathematically defined by: $CVaR_\alpha(X) = \mathbb{E}[X \mid X \geq VaR_\alpha(X)]$ where X is a real random variable. The benefit of using CVaR instead of VaR is that CVaR takes into account not only the VaR value but also the tail information of a distribution. CVaR is known to satisfy FSD [3].

Example 3. Expected Utility (EU) is a well-known decision criterion in decision under risk [20] and decision under uncertainty [29], which is known to satisfy FSD [3]. It is defined as follows: $EU(X) = \mathbb{E}(u(X))$ where X is a real random variable and $u : \mathbb{R} \to \mathbb{R}$ is a so-called von Neumann-Morgenstern utility function. The utility value $u(x)$ represents how much x is valuable. For this reason, function u is assumed to be monotonic (i.e., in our settings, $x \leq y \Rightarrow u(x) \geq u(y)$). In decision theory, it is well-known that a concave (resp. convex) function u leads to a risk-averse (resp. risk-seeking) decision criterion. Although we focus on risk-averse criteria in this paper (as it is most people's concern in transportation), note that our approach could also tackle the risk-seeking case.

There are many other possible examples of ρ that satisfies property FSD: for instance, semideviations [22], rank-dependent utility [27], Yaari's dual model [32]... Our solution algorithm covers all those cases.

5 Solution Algorithm

We propose an algorithm that is an adapted version of the standard A* algorithm to solve the proposed risk-averse shortest path problem using time-dependent stochastic costs. It generalizes the algorithm proposed by Chen et al. [6] to general ρ measures that satisfies FSD and extends the algorithm proposed by Parmentier et al. [24] to the time-dependent cost setting.

The proposed algorithm keeps the basic features of the standard A* algorithm, for example, an open set $\mathcal{O}$ while adding new features such as labeling with random variables and path pruning using FSD dominance. We now explain why the notions of label (for evaluating the value of a subpath ending in node n) and priority (for guiding the order the subpaths are examined) need to be redefined in our setting and how they can be redefined. In the standard A* algorithm for computing a shortest path, a node n in the priority queue $\mathcal{O}$ is the end node of a subpath for which only one label (i.e., cumulated cost $g(n) = c(\pi_{on})$) needs to be stored. This is possible because we have $\forall n \in V$:

$$c(\pi_{on}) \leq c(\pi'_{on}) \implies c(\pi_{on} \oplus \pi_{nd}) \leq c(\pi'_{on} \oplus \pi_{nd})$$

where π_{on} and π'_{on} are two paths from node o to node n and π_{nd} is a path from n to d. Unfortunately, in our setting, a counterpart of these inequalities with respect to ρ does not hold, due to the possible non-linearity of criterion ρ:

$$\rho(C_t(\pi_{on})) \leq \rho(C_t(\pi'_{on})) \not\Longrightarrow \rho(C_t(\pi_{on} \oplus \pi_{nd})) \leq \rho(C_t(\pi'_{on} \oplus \pi_{nd})) \quad (5)$$

In words, a dominated subpath can become non-dominated when extended.

Example 4. We give an example for the case when ρ is VaR with $\alpha = 95\%$. Assume the probability distributions are given in Table 1. One can check that:

$$VaR(C_t(\pi_{on})) = 1 < 2 = VaR(C_t(\pi'_{on})) \text{ and}$$
$$VaR(C_t(\pi_{on} \oplus \pi_{nd})) = 3 > 2 = VaR(C_t(\pi'_{on} \oplus \pi_{nd}))$$

As a consequence of (5), labels have a more complex form. Following previous

Algorithm 1: Proposed adapted A* algorithm

Data: graph $G = (V, E)$, random costs C_t, heuristic h, criterion ρ, upperbound UB, origin node o, destination node d, departure time t

Result: risk-averse shortest path

```
begin
    O ← {(o, 0)}
    while O ≠ ∅ do
        (n, C) ← highest priority pair in O
        if n = d then return corresponding path remove (n, C) from O
        for n' ∈ E+(n) do
            C' ← C + C_{t+C}(n, n')        ▷initial time t + C selects the edge cost
            f(n', C') ← ρ(C' + h(n'))
            if n' = d and f(n', C') < UB then
                UB ← f(n', C')
            else
                if f(n', C') ≥ UB then continue
            if n' ∉ O then
                add (n', C') in O
            else
                if C' not FSD-dominating any (n', C'') ∈ O then
                    add (n', C') in O and remove FSD-dominating (n', C'') ∈ O
                else
                    continue
    return ∅
```

Table 1: Cumulative distributions.

x	0	1	2	3	4
$F_{C_t(\pi_{on})}$	0	0.95	1	1	1
$F_{C_t(\pi'_{on})}$	0.9	0.9	1	1	1
$F_{C_t(\pi_{nd})}$	0.8	0.9	1	1	1
$F_{C_t(\pi_{on} \oplus \pi_{nd})}$	0	0.76	0.895	0.995	1
$F_{C_t(\pi'_{on} \oplus \pi_{nd})}$	0.72	0.81	0.98	0.99	1

related work [21, 6, 24], the label of node n is defined as $C_t(\pi_{on})$ instead of $\rho(C_t(\pi_{on}))$. For a given node, two labels can be compared with FSD-dominance (thanks to Corollary 5.2). As it is a partial order, a node can then receive several labels. For this reason, elements of $\mathcal{O}$ are pairs (n, C) where n is a node and C is a random variable representing the cost of a subpath from node o to node n.

The priority of a pair (n, C) in $\mathcal{O}$ is defined as $f(n, C) = \rho(C + h(n))$ where $h(n)$ is a known heuristic evaluation of a subpath from node n to node o. We assume that $h(n)$ is FSD-dominated by the random cost of any subpath from node n to node d. Heuristic $h(n)$ can be a deterministic value [6] as usual or

more generally a random variable [24].

Defining the label as such and comparing them with FSD dominance are justified because of the following lemma [21, 6, 24] and corollary, written for X, Y, Z three real random variables.

Lemma 5.1 *If $F_X \succsim_{FSD} F_Y$, then $F_{X+Z} \succsim_{FSD} F_{Y+Z}$.*

This lemma can be interpreted in our context as follows: If the label (i.e., random variable or its associated probability distribution more exactly) of a subpath π ending in n FSD-dominates the label of another subpath also ending in n, then any extension of those two subpaths will keep the direction of the dominance.

As a direct consequence of Lemma 5.1 and FSD, we have:

Corollary 5.2 *If $F_X \succsim_{FSD} F_Y$, then $\rho(X + Z) \geq \rho(Y + Z)$.*

This corollary states that if a given node n has two labels, one FSD-dominating the other, the former label can be pruned as it will lead to a higher ρ value.

Thanks to INC, the following proposition explains why it is sound to end the algorithm as soon as node d is examined (Line 5 of Algorithm 1).

Proposition 5.3 *When node d is chosen, the corresponding path is ρ-minimum.*

Property INC was not considered in Parmentier et al.'s work [24]. Contrary to their algorithm, ours can stop as soon as a path to node d is found.

In order to avoid generating too many subpaths, we use an upperbound UB on the best ρ value known so far. When starting Algorithm 1, we can use $UB = +\infty$ or better compute a standard shortest path and use its ρ value as an upperbound. Then, UB can be updated each time a path to d is found (Line 9). Besides, this algorithm can be sped up by pruning with any other known lower bound to the ρ value (see the case study where we use the expected duration).

Note that in general, Line 5 may be hard to compute. In our case study, we assume the time is discretized into equal-length intervals on which probability distributions are assumed to be constant. Moreover, we also assume all distributions are discretized. In the next section, we explain this in more details.

6 Case Study

We demonstrate our algorithm with ρ chosen as the conditional value-at-risk (CVaR) with $\alpha = 90\%$. This seems to be a better choice than VaR, which was used in Chen et al.'s work [6], because it not only takes into account the VaR threshold, but also the tail distribution. Besides, being a coherent risk measure [1], it enjoys nicer properties than VaR. We implemented our adapted A* algorithm in OpenTripPlanner[3], an open-source platform for route planning, which offers a map-based web interface and standard shortest path algorithms.

In order to work with real traffic data, we estimated the dynamic random costs from taxi trip data[4] released by the New York City TLC (Taxi and Limousine Commission). We first explain how a probability distribution for the random duration of an edge was estimated and then present an illustration of results that can be obtained thanks to our algorithm.

[3] `http://www.opentripplanner.org`

[4] `http://www.nyc.gov/html/tlc/html/about/trip_record_data.shtml`

6.0.1 Data Cleaning and Estimation.

The dataset contains records of taxi trips in Manhattan from 2009 to 2015. We only used the data from 2011 to 2015, as the data size was large and we preferred focusing on the most recent records. The dataset contains trip information including pick-up/drop-off locations and pick-up/drop-off times. We only took into account trips inside the Manhattan area, which represents a network of 5,111 nodes and 16,396 edges. During the data cleaning phase, we filtered out trips that had a pick-up or drop-off location outside Manhattan. We also removed abnormal trips, which may be due to incorrect GPS readings.

Because the actual path of a trip and detailed times at each intersection of a trip were not provided, we had to make two assumptions to extract random duration $C_t(e)$ of an edge $e \in E$ from the dataset:

A1 A trip follows the shortest path from origin to destination.

A2 The driver maintains the same speed along the trip.

Given the nature of the dataset, the assumptions seem reasonable enough. A1 leads to a small overestimation of travel durations in each edge. For our risk-averse route planning problem, overestimation is better than underestimation. A2 is a simplifying assumption, which neglects the effects of traffic lights, intersections, turns... We do not think it has a too big impact for our application, especially given that we have already overestimated the durations.

Based on A1, for each trip, we computed its shortest path from its origin to its destination using standard A* in terms of duration, where the duration of an edge (i.e., portion of a street) equals to the length (i.e., distance) of an edge divided by the maximum speed limit allowed in that edge. Then, given the computed shortest path π, we could generate a duration sample for each of its edge based on A2 with $c_e = c_\pi \times \frac{l_e}{l_\pi}$ where c_e is the duration of an edge e of π, c_π the total duration of the trip, l_e the length of edge e and l_π the length of π.

Samples c_e's were then collected and used to estimate $P_{C_t(e)}$. As we expect different traffic patterns on weekdays and during weekends, we divided the days of a week into two classes: $Weekdays = \{Mon., Tues., Wed., Thur., Fri.\}$ and $Weekends = \{Sat., Sun.\}$. We divided a day into 24 bins of 1 hour. For a specific edge e, we obtained 24 distributions $P_{C_t(e)}$ (one for each hour) and assume the distribution was constant during an interval of one hour. Moreover, we assume those distributions are discrete and defined over 100 bins of 6 seconds. Durations that exceeds 600 seconds were counted as 600 seconds.

6.0.2 Experimental Results.

With the adapted A* algorithm described before, we can find the risk-averse path between any pair of origin and destination. For this case study, following Chen et al.'s work [6], we define the heuristic function used in the implemented risk-averse path finding system as $h(n) = \frac{d(n)}{v_{max}}$ where $d(n)$ is the shortest length of a path from node n to node d and v_{max} is the maximum travel speed in the network. Therefore, $h(n)$ is the shortest possible duration to go from n to d.

Besides, it is known that $CVaR_\alpha$ is increasing with α and $CVaR_{0\%}$ is the expectation. Therefore, we also maintained an expected duration of a subpath π_{on} and estimated a lowerbound of the expected duration of an extension of π_{on} to node d (as in standard A*). This lower bound can be used to prune subpaths by comparing it with the upperbound UB.

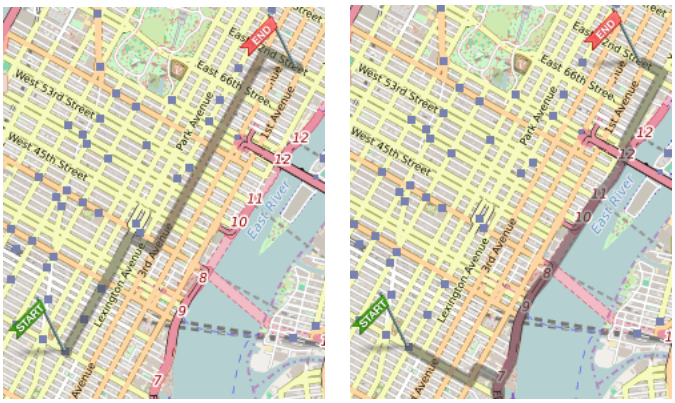

(a) path at 6:00 a.m. on Wednesday

(b) path at 8:00 a.m. on Wednesday

Figure 2: Examples of risk-averse paths

To illustrate our system, we present one example where for the same pair of origin and destination nodes, CVaR yields different risk-averse paths depending on the departure time. As depicted in Figure 2, at 6:00 a.m. on Wednesday, before rush hours, the risk-averse path (with a CVaR of 24 mins) is very similar to the shortest distance path because there is little risk of congestion. Its CVaR can be interpreted as follows: In the worst 10% of the case, the average duration of the trip will be 24 minutes. And, in most cases, the observed travel duration would be much less than 24 minutes. In contrast, at 8:00 a.m. on the same day during rush hour, the risk-averse path is no longer the shortest distance path, but a path (with a CVaR of 31 mins) that passes via a highway, which has less probability of congestion. Although the risk-averse path may be a longer path to drive, it is a less risky path in terms of CVaR.

The computation times depend on the origin and destination nodes. By averaging over 100 runs where those pairs where selected randomly, the average computation time was less than one second (976.2 millisecs) using a computer equipped with an *Intel Xeon E31225 @ 3.10GHz*. To make this system usable in a real application, the computation time could be further improved. We expect this could be achieved with different optimization techniques: e.g., memoization, better heuristics, fitting duration samples to continuous distributions... As we wanted to demonstrate the feasibility of our approach, we leave this as future work.

7 Conclusion

In this paper, we proposed an adapted A* algorithm, which accommodates any risk measure or decision criterion that is monotonic with first-order stochastic dominance, to find a risk-averse shortest path in a transportation network with time-dependent stochastic costs. Besides, we demonstrated our algorithm on a case study with NYC taxi data and obtained reasonable results.

As future work, we plan to improve the computational efficiency of our method, taking inspiration from the techniques developed for standard shortest path problems [2, 7]. Moreover, we would like to test our system on more

accurate historical traffic data. Finally, we plan to extend the approach to take into account other kinds of costs, such as power consumption.

References

[1] Artzner, P., Delbaen, F., Eber, J., Heath, D.: Coherent measures of risk. Mathematical Finance 9(3), 203–228 (1999)

[2] Bast, H., Delling, D., Goldberg, A., Müller-Hannemann, M., Pajor, T., Sanders, P., Wagner, D., Werneck, R.: Route planning in transportation networks (2015), arXiv:1504.05140v1

[3] Bäuerle, N., Müller, A.: Stochastic orders and risk measures: Consistency and bounds. Mathematics & Economics (2005)

[4] Bellman, R.: On a routing problem. Quarterly of Applied Mathematics 16, 87–90 (1958)

[5] Bertsekas, D., Tsitsiklis, J.: An analysis of stochastic shortest paths problems. Mathematics of Operations Research 16, 580–595 (1991)

[6] Chen, B.Y., Lam, W.H.K., Sumalee, A., Li, Q., Tam, M.L.: Reliable shortest path problems in stochastic time-dependent networks. Journal of Intelligent Transportation Systems 18(2), 177–189 (2014)

[7] Delling, D., Sanders, P., Schultes, D., Wagner, D.: Engineering route planning algorithms. Algorithmics, LNCS 5515, 117–139 (2009)

[8] Dijkstra, E.: A note on two problems in connexion with graphs. Numerische Mathematik 1, 269–271 (1959)

[9] Dreyfus, S.: An appraisal of some shortest-path algorithms. Operations Research 17(3), 395–412 (1969)

[10] Embrechts, P., Kluppelberg, C., Mikosch, T.: Modelling Extremal Events for Insurance and Finance. Springer (1997)

[11] Ford, L.J.: Network flow theory. Tech. rep., Rand Corporation (1956)

[12] Frank, H.: Shortest paths in probabilistic graphs. Operations Research 17(4), 583–599 (1969)

[13] Fu, L., Rilett, L.: Expected shortest paths in dynamic and stochastic traffic networks. Transportation Research Part B: Methodological 32(7), 499–516 (1998)

[14] Gavriel, C., Hanasusanto, G., Kuhn, D.: Risk-averse shortest path problems. In: IEEE 51st Annual Conference on Decision and Control. pp. 2533–2538 (2012)

[15] Goldberg, A., Harrelson, C.: Computing the shortest path: A* meets graph theory. In: SODA. pp. 156–165 (2005)

[16] Hart, P.E., Nilsson, N.J., Raphael, B.: A formal basis for the heuristic determination of minimum cost paths. IEEE Trans. Syst. and Cyb. 4(2), 100–107 (1968)

[17] Jorion, P.: Value-at-Risk: The New Benchmark for Managing Financial Risk. McGraw-Hill (2006)

[18] Kaufman, D., Smith, R.: Fastest paths in time-dependent networks for intelligent vehicle-highway systems application. Journal of Intelligent Transportation Systems 1(1), 1–11 (1993)

[19] Moore, E.F.: The shortest path through a maze. In: Proc. of the Int. Symp. on the Theory of Switching. pp. 285–292 (1959)

[20] von Neumann, J., Morgenstern, O.: Theory of games and economic behavior. Princeton university press (1944)

[21] Nie, Y., Wu, X.: Shortest path problem considering on-time arrival probability. Transportation Research Part B: Methodological 43(6), 597–613 (2009)

[22] Ogryczak, W., Ruszczynski, A.: From stochastic dominance to mean-risk models: Semideviations as risk measures. European Journal of Operational Research 116(33–50) (1999)

[23] Orda, A., Rom, R.: Shortest-path and minimum delay algorithms in networks with time-dependent edge-length. Journal of the ACM 37(3), 607–625 (1990)

[24] Parmentier, A., Meunier, F.: Stochastic shortest paths and risk measures. In: arXiv preprint (2014)

[25] Peyer, S., RautenBach, D., Vygen, J.: A generalization of dijkstra's shortest path algorithm with applications to vlsi routing. Journal of Discrete Algorithms 7(4), 377–390 (2009)

[26] Puterman, M.: Markov decision processes: discrete stochastic dynamic programming. Wiley (1994)

[27] Quiggin, J.: Generalized expected utility theory: the rank-dependent model. Kluwer Academic Publishers (1993)

[28] Russell, S., Norvig, P.: Artificial Intelligence: A Modern Approach. Prentice-Hall, 2nd edn. (2003)

[29] Savage, L.: The foundations of statistics. J. Wiley and sons (1954)

[30] Shaked, M., Shanthikumar, J.: Stochastic Orders and Their Applications. Academic press (1994)

[31] Sigal, C., Pritsker, A., Solberg, J.: The stochastic shortest route problem. Operations Research 28, 1122–1129 (1980)

[32] Yaari, M.: The dual theory of choice under risk. Econometrica 55, 95–115 (1987)